



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
12.10.2016 Bulletin 2016/41

(51) Int Cl.:
H04L 29/06 (2006.01)

(21) Application number: **16164616.1**

(22) Date of filing: **11.04.2016**

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR
Designated Extension States:
BA ME
Designated Validation States:
MA MD

(72) Inventors:
• **Hassanzadeh, Amin**
Arlington, VA Virginia 22203 (US)
• **Modi, Shimon**
Washington, DC District of Columbia 20001 (US)
• **Mulchandani, Shaan**
Arlington, VA Virginia 22209 (US)
• **Negm, Walid**
Reston, VA Virginia 20190 (US)

(30) Priority: **09.04.2015 US 201562145248 P**
31.08.2015 US 201514841227

(74) Representative: **Conroy, John**
Fish & Richardson P.C.
Highlight Business Towers
Mies-van-der-Rohe-Straße 8
80807 München (DE)

(71) Applicant: **Accenture Global Services Limited**
Dublin 4 (IE)

(54) **EVENT CORRELATION ACROSS HETEROGENEOUS OPERATIONS**

(57) Systems, methods, and apparatus, including computer programs encoded on computer storage media, for performing event correlation across heterogeneous operations.




FIG. 1

Description

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefit of U.S. Provisional Application No. 62/145,248, filed April 9, 2015, and titled "Event Correlation Across Heterogeneous Operations," which is incorporated by reference.

BACKGROUND

[0002] The present disclosure relates to security and network operations.

SUMMARY

[0003] In general, one innovative aspect of the subject matter described in this specification can be embodied in methods for correlating domain activity, including receiving first domain activity data from a first network domain and second domain activity data from a second network domain, the first domain activity data and the second domain activity data including events, alerts, or both from the respective first and second network domains, filtering the first domain activity data and the second domain activity data to remove irrelevant activity data, based on a first set of profile data for devices in the first network domain and a second set of profile data for devices in the second network domain, aggregating unfiltered first domain activity data and unfiltered second domain activity data, correlating aggregated unfiltered first domain activity data and unfiltered second domain activity data to determine an attack path for an attack that occurs across the first network domain and the second network domain, based on attack signatures and profiles associated with previously identified attacks, and generating a visualization of the attack path.

[0004] Other embodiments of this aspect include corresponding computer methods, and include corresponding apparatus and computer programs recorded on one or more computer storage devices, each configured to perform the actions of the methods. A system of one or more computers can be configured to perform particular operations or actions by virtue of having software, firmware, hardware, or a combination of them installed on the system that in operation causes or cause the system to perform the actions. One or more computer programs can be configured to perform particular operations or actions by virtue of including instructions that, when executed by data processing apparatus, cause the apparatus to perform the actions.

[0005] These and other embodiments may each optionally include one or more of the following features. For instance, the first domain activity data, the second domain activity data, or both, can include log data provided by one or more security sensors. The first network domain can be an information technology network domain and the second network domain can be an operational

technology network domain. The filtering can include determining, for each event or alert, a corresponding attack and a corresponding target, determining, based on profile data for the corresponding target, that the attack on the target is rendered unsuccessful, and filtering the corresponding event or alert. The filtering can include, for each unfiltered event or alert, dynamically retrieving current status information about the corresponding target, determining, based the current status information about the corresponding target, that the attack on the target is rendered unsuccessful, and filtering the corresponding event or alert.

[0006] Particular embodiments of the subject matter described in this specification may be implemented so as to realize one or more of the following advantages. A single vantage point and a standardized data exchange format may be provided for analyzing event/alert log data from information technology (IT) and operational technology (OT) networks. Currently available security sensors (e.g., intrusion detection systems (IDS), intrusion prevention systems (IPS), and other suitable security sensors) may be leveraged, resulting in architectural independence, flexibility, and compatibility with legacy infrastructures/networks. False alarm rates may be reduced in an event detection process. Multi-step, multi-domain threat scenarios may be detected and/or constructed. Complex scenarios may be visually represented, and network status may be shown at each step of an attack. Threat intelligence platforms/services may be integrated to further enrich and contextualize information and constructed scenarios.

[0007] The details of one or more implementations of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other potential features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

DESCRIPTION OF DRAWINGS

[0008]

FIGS. 1 - 3 depict example systems that can execute implementations of the present disclosure.

FIG. 4A depicts an example system that can execute implementations of the present disclosure.

FIG. 4B depicts an example data structure that can be used by implementations of the present disclosure.

FIG. 5A depicts an example system that can execute implementations of the present disclosure.

FIG. 5B and 5C depict example data structures that can be used by implementations of the present disclosure.

FIG. 6 is a flowchart of an example process that can be executed in accordance with implementations of the present disclosure.

FIG. 7 is a block diagram of a computing system that can be used in connection with computer-implemented methods described in this document.

[0009] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0010] This specification describes systems, methods, and computer programs for performing event correlation across heterogeneous operations. For example, an industrial internet may be used to manage and administer industrial control systems (ICS), which may communicate over an enterprise network and may include information technology (IT) and operational technology (OT) domains. Some threat scenarios may include multi-step, multi-domain attacks, and may include attacks that originate in one domain, and proceed to another domain. By filtering, aggregating, and correlating data from event/alert logs from each domain (e.g., IT and OT domains), for example, complex attack patterns may be detected. Information about the attack patterns (e.g., visualization data) may be reported to a security analyst, and may be used for implementing appropriate courses of action.

[0011] FIG. 1 depicts an example system 100 that can execute implementations of the present disclosure. In the present example, the system 100 (e.g., an industrial control system) includes multiple network domains, including an information technology (IT) network domain 102 (e.g., including an enterprise network) and an operational technology (OT) network domain 104. The information technology network domain 102 and the operational technology network domain 104 can be in communication, for example, over a demilitarized zone (DMZ) 106a of the information technology network domain 102 and a demilitarized zone (DMZ) 106b of the operational technology network domain 104. Each of the network domains 102 and 104, for example, may include local and wide area networks (LAN/WAN) and wireless networks, and can be used to integrate various computing devices, such as servers, mainframes, desktops, laptops, tablets, smartphones, and industrial control devices and sensors, that may run on multiple different operating systems and may employ multiple different communication protocols.

[0012] The information technology network domain 102 can include various computing devices (e.g., computing server 112), input/output devices (e.g., interface device 114), and/or subsystems. The computing server 112, for example, can include one or more processors configured to execute instructions stored by computer-readable media for performing various operations, such

as input/output, communication, data processing, and/or data maintenance. To interact with the computing server, for example, a user can employ the interface device 114 (e.g., including one or more presentation components, such as a display, and one or more input components such as a keyboard, mouse, and/or touchpad).

[0013] The operational technology network domain 104 can include various computing devices, input/output devices, and/or subsystems. In the present example, the operational technology network domain 104 includes a supervisory system 120, a historian server 122, an application server 124, one or more human-machine interface (HMI) devices (e.g., HMI device 126), and one or more controller devices (e.g., controller device 128) and sensor devices (e.g., sensor device 130). The supervisory system 120, for example, can coordinate one or more low-level controls and/or low-level sensors. In the present example, the supervisory system 120 can provide data to and receive data from the controller device 128 and the sensor device 130. The historian server 122, for example, can store, maintain, and provide information related to activities performed by each controller device and sensor data provided by each sensor device in the operational technology network domain 104. The application server 124, for example, can host applications that may operate within the operational technology network domain 104.

[0014] In some implementations, the system 100 may include one or more security sensors (e.g., security sensors 108a, 108b, 108c, and 108d). In general, security sensors included in the system 100 may include network based (NIDS) and host based (HIDS) intrusion detection systems, intrusion prevention systems (IPS), anti-virus systems, firewalls, and other detection/logging services (e.g., web server logs, database logs, etc.) which can monitor communications activity to and from computing devices included in the industrial technology (IT) network domain 102, the IT DMZ 106a, the operational technology (OT) network domain 104, and/or the OT DMZ 106b, and can monitor system activity associated with the devices. Data associated with potentially malicious activity may be detected (and optionally, recorded) by the security sensors 108a, 108b, 108c, and 108d (e.g., as event/alert data, log files, etc.), and/or other detection/logging devices included in the system 100, and/or may be provided to other components of the system 100. For example, activity data 140a, 140b, 140c, and 140d (e.g., detected by the corresponding security sensors 108a, 108b, 108c, and 108d) may be provided to an event correlation system 150. Such activity data may also be provided to an event correlation system 150 by a Security Information and Event Management (SIEM) system. The activity data 140a, for example, may include enterprise data from the information technology network domain 102, provided by host-based monitoring systems (e.g., intrusion detection/prevention systems, web server logging services, system logs, etc.) and/or network-based monitoring systems (e.g., intrusion detection/prevention

systems, firewalls, routers, etc.). The activity data 140b, for example, may include data associated with communication over the IT DMZ 106a. The activity data 140c, for example, may include data associated with communication over the OT DMZ 106b. The activity data 140d, for example, may include supervisory data, control layer data, and/or sensor and controller device data from the operational technology network domain 104, provided by host-based monitoring systems and/or network-based monitoring systems.

[0015] In the present example, each of the activity data 140a, 140b, 140c, and 140d may include event and/or alert data. In general, events are atomic pieces of data associated with communications and system activity, whereas alerts may be triggered in response to an event or a sequence of events. Data provided by the security sensors 108a, 108b, 108c, and 108d, for example, may include alert data. Data provided by a host (e.g., the computing server 112), controller device (e.g., the controller device 128) or sensor device (e.g., the sensor device 130), or data included in log files, for example, may include event data.

[0016] The event correlation system 150, for example, can receive the activity data 140a, 140b, 140c, and 140d from multiple domains (e.g., the information technology (IT) network domain 102, the IT DMZ 106a, the operational technology (OT) network domain 104, and the OT DMZ 106b), and can standardize, filter, aggregate, and correlate the data to detect anomalies and potentially malicious activity associated with multi-stage, multi-domain attacks. In the present example, the event correlation system 150 can include various computing devices (e.g., computing server 152), input/output devices (e.g., interface device 154), and/or subsystems. The computing server 152, for example, can include one or more processors configured to execute instructions stored by computer-readable media for performing various operations, such as input/output, communication, data processing, and/or data maintenance. To interact with the computing server, for example, a user can employ the interface device 154 (e.g., including one or more presentation components, such as a display, and one or more input components such as a keyboard, mouse, and/or touchpad).

[0017] In some implementations, output may be provided by the event correlation system 150 to another system (e.g., a security information and event management (SIEM) system) and/or to a system operator as reporting/visualization data. Based on the system output, for example, appropriate courses of action may be employed to counter ongoing and/or future attacks. In the present example, the information technology (IT) network domain 102, the IT DMZ 106a, the operational technology (OT) network domain 104, and the OT DMZ 106b each has different characteristics (e.g., architecture, resources, protocols, and standards), and each domain may be susceptible to different security threats. Occasionally, correlations may not be detected among events/alerts within a single domain, (and if correlations are detected, an ex-

tent of an associated compromise may not be entirely known), but correlations may be detected among events/alerts across multiple domains. By correlating data from multiple domains, for example, complex attacks (e.g., multistage, multi-domain attacks executed over time) may be detected, and a single vantage point may be provided to security technicians.

[0018] FIG. 2 depicts an example system 200 that can execute implementations of the present disclosure. In the present example, the system 200 includes an event correlation system 202 (e.g., similar to the event correlation system 150, shown in FIG. 1). The event correlation system 202, for example, can include various hardware and/or software-based components (e.g., software modules, objects, engines, libraries, etc.) including an information technology (IT) activity data filter 210, an operational technology (OT) activity data filter 212, an aggregator 214, a correlator 216, a pattern recognizer 218, a response generator 220, and an operator input/output (I/O) component 222. Various data sources (e.g., databases, file systems, etc.) may maintain data used by the system 200 and its components. In the present example, the system 200 can receive information from an information technology (IT) network data source 230, an operational technology (OT) network data source 232, and a threat intelligence data source 234. Activity data associated with a demilitarized zone (DMZ) or peripheral network, for example, may be provided by the information technology network data source 230 and/or the operational technology network data source 232. In general, the system 200 and its various components (e.g., components 210, 212, 214, 216, 218, 220, and 222) can perform functions for processing event/alert data received from various sources, aggregating the data, correlating the data, detecting patterns in the data, and providing relevant information to system operators and/or other systems.

[0019] In the present example, the event correlation system 202 can receive information technology (IT) activity data 240 that includes event/alert data from an information technology network (e.g., the information technology (IT) network domain 102, and optionally the IT DMZ 106a, shown in FIG. 1), and can receive operational technology (OT) activity data 242 that includes event/alert data from an operational technology network (e.g., the operational technology (OT) network domain 104, and optionally the OT DMZ 106b, shown in FIG. 1). In some implementations, the information technology activity data 240 and/or the operational technology activity data 242 may include log data provided by one or more security sensors (e.g., the security sensors 108a, 108b, 108c, and 108d, shown in FIG. 1). Upon receiving the information technology activity data 240, for example, the event correlation system 202 can use the information technology activity data filter 210 to filter out irrelevant (or "false") events/alerts, based on data provided by the information technology network data source 230. Similarly, upon receiving the operational technology activity

data 242, for example, the event correlation system 202 can use the operational technology activity data filter 212 to filter out irrelevant (or "false") events/alerts, based on data provided by the operational technology network data source 232. Operation of the information technology activity data filter 210 and the operational technology activity data filter 212 is discussed in further detail below in association with FIG. 3.

[0020] After filtering the information technology activity data 240 and the operational technology activity data 242, for example, filtered event/alert data can be provided by the information technology activity data filter 210 and the operational technology activity data filter 212 to the aggregator 214. In general, the event correlation system 202 can use the aggregator 214 to remove duplicate and/or redundant events/alerts, to combine events/alerts related to the same attack, and to combine events/alerts relating to different attacks but possessing similar characteristics, thus reducing the number of events/alerts under consideration. In some implementations, the aggregator 214 may reference data provided by the information technology network data source 230 and/or the operational technology network data source 232 when performing aggregation operations. Operation of the aggregator 214 is discussed in further detail below in association with FIG. 4A.

[0021] After aggregating the event/alert data, for example, aggregated data can be provided by the aggregator 214 to the correlator 216. In general, the event correlation system 202 can use the correlator 216 to generate a chain of events/alerts that may correspond to a threat scenario, and the event correlation system 202 can use the pattern recognizer 218 (e.g., based on data provided by the threat intelligence data source 234) to identify attack patterns associated with the threat scenario, and to further describe and/or enrich threat scenario information. Based on threat scenarios identified by the correlator 216 and attack patterns identified by the pattern recognizer 218, and optionally based on operator input received by the operator input/output component 222, for example, the response generator 220 can provide appropriate courses of action for responding to threats to the information technology network 250 and the operational technology network 252. Operation of the correlator 216 and the pattern recognizer 218 is discussed in further detail below in association with FIG. 5A.

[0022] FIG. 3 depicts an example system 300 that can execute implementations of the present disclosure. In the present example, the system 300 includes a filtering and verification system 302 (e.g., corresponding to the information technology activity data filter 210 and the operational technology activity data filter 212, shown in FIG. 2). The filtering and verification system 302, for example, can include various hardware and/or software-based components (e.g., software modules, objects, engines, libraries, etc.) including a rule-based filter 310 for information technology (IT) activity data, a rule-based filter 312 for operational technology (OT) activity data, an op-

tional verifier 320 for information technology (IT) activity data, and an optional verifier 322 for operational technology (OT) activity data. Various data sources (e.g., databases, file systems, etc.) may maintain data used by the system 300 and its components. In the present example, the system 300 includes an information technology (IT) network data source 330 (e.g., including configuration management information associated with devices in the information technology (IT) network domain 102, shown in FIG. 1), an operational technology (OT) network data source 332 (e.g., including configuration management information associated with devices in the operational technology (OT) network domain 104, shown in FIG. 1), and a filtered alerts data source 336. In general, the system 300 and its various components (e.g., components 310, 312, 320, and 322) can perform functions for processing event/alert data received from various different sources. By removing or filtering out irrelevant event/alert data (i.e., false positives and/or noise), for example, the accuracy of correlation engines may be increased.

[0023] In the present example, the filtering and verification system 302 can receive information technology (IT) activity data 340 that includes event/alert data from an information technology network, and optionally, a corresponding DMZ (e.g., the information technology (IT) network domain 102 and the IT DMZ 106a, shown in FIG. 1), and can receive operational technology (OT) activity data 342 that includes event/alert data from an operational technology network, and optionally, a corresponding DMZ (e.g., the operational technology (OT) network domain 104 and the DMZ 106b, shown in FIG. 1). In some implementations, the information technology activity data 340 and/or the operational technology activity data 342 may include log data provided by one or more security sensors (e.g., security sensors 108a, 108b, 108c, and 108d, shown in FIG. 1). For example, activity data received from multiple sources (e.g., multiple security sensors, intrusion detection systems, and/or other security tools) may be heterogeneous in regard to language, protocols, and standards. Such activity data may be heterogeneous, for example, not only because of different security tools in a single domain (which may be resolved through the use of alert/event standardization/normalization tools that convert data to a standard format), but because of different protocol standards which may be employed in multiple different domains by the same security tool. As another example, a standard format may be used for communicating activity data. Upon receiving the information technology activity data 340, for example, the filtering and verification system 302 can use the rule-based filter 310 for information technology activity data to filter out irrelevant (or "false") events/alerts, based on data provided by the information technology network data source 330 (e.g., similar to the information technology network data source 230, shown in FIG. 2), and in conjunction with additional rules that may be defined by system administrators. Similarly, upon receiving the opera-

tional technology activity data 342, for example, the filtering and verification system 302 can use the rule-based filter 312 for operational technology activity data to filter out irrelevant (or "false") events/alerts, based on data provided by the operational technology network database 332 (e.g., similar to the operational technology network data source 232, shown in FIG. 2), and in conjunction with additional rules that may be defined by system administrators.

[0024] In general, rule-based filtering performed by each of the rule-based filters 310 and 312 can remove irrelevant events/alerts (e.g., events/alerts that are not determined to be associated with a potential attack) based on a target's profile and/or characteristics of the events/alerts. Rule-based filtering, for example, may apply to defined rules that discard particular events/alerts (e.g., false positives) based on how frequently events/alerts with certain characteristics occur, and their relative rate of change with regard to occurrence. Profile data for potential targets (e.g., computing devices) in the information technology network domain 102 (shown in FIG. 1) can be maintained by the information technology network data source 330, and profile data for potential targets (e.g., computing devices, controllers, and sensors) in the operational technology network domain 104 (shown in FIG. 1) can be maintained by the operational technology network data source 332. For each received event/alert, for example, an appropriate rule-based filter may reference profile data from an appropriate network data source for a target that corresponds to the event/alert (e.g., based on device address), and can determine whether the received event/alert indicates a potential attack. For example, a network-based intrusion detection system may not have specific information about an attacker or about a target, but may generate an alert based on the contents of a communications packet - that is, the alert may be generated if the packet includes an exploit directed to a known vulnerability. However, the generated alert in the present example may or may not indicate a successful attack on the target. For example, if an attack relies on certain system attributes (e.g., a type of operating system), but the system has different attributes (e.g., a different operating system) that are not affected by an attempted attack, the attack is rendered unsuccessful. As another example, if a communications packet is directed to a computing device that does not exist on a network, the network drops the communications packet, rendering the packet ineffective and the attack unsuccessful. By filtering events/alerts associated with attack attempts that are likely to be unsuccessful, reconnaissance attempts (intentional or unintentional), and/or internal activity known to be benign, for example, the number of events/alerts under consideration may be reduced, thus reducing the amount of processing in subsequent stages.

[0025] In some implementations, profile data and/or statuses of potential targets (e.g., computing devices) may be dynamically determined when filtering received

events/alerts. For example, after performing rule-based filtering on the information technology activity data 340, the filtering and verification system 302 can optionally use the verifier 320 to dynamically verify profile data for a target on the information technology network domain 102 (shown in FIG. 1), and the filtering and verification system 302 can optionally use the verifier 322 to dynamically verify profile data for a target on the operational technology network domain 104 (shown in FIG. 1). For each received event/alert that has not been previously filtered, for example, an appropriate verifier can determine whether dynamically retrieved information regarding a target corresponding to the event/alert (e.g., based on device address) indicates that the event/alert is associated with actual malicious activity. For example, for an unfiltered event/alert corresponding to a target on the information technology network domain 102, the verifier 320 can communicate with a system/network scanner 350 (e.g., with access to configuration management information) to receive current information about the target. In the present example, a communications packet may be directed to an open port on the target, and the rule-based filter 310 may not filter the corresponding alert. However, based on information provided by the system/network scanner 350, the verifier 320 may determine that the target has been patched to counter an attack associated with the alert, for example, rendering the communications packet ineffective and the attack unsuccessful. In the present example, the alert may be filtered out and/or labeled as a potentially unsuccessful attack. As another example, for an unfiltered event/alert corresponding to a target on the operational technology network domain 104, the verifier 322 can communicate with a device virtualization component 352 to receive current information about the target. In the present example, a communications packet may be directed to changing a setting (e.g., a temperature setting, an on/off setting, a power level, a position, etc.) on the target (e.g., a controller device). The device virtualization component 352, for example, can query the target (or one or more sensors associated with the target) for its status to determine an effect of the communications packet. A negative effect, for example, may indicate a potentially successful attack, whereas a neutral effect or lack of an effect may indicate a potentially unsuccessful attack.

[0026] After performing rule-based filtering and verification, for example, the filtering and verification system 302 can record filtered alert data associated with potentially unsuccessful attacks and/or false positives (e.g., in the filtered alerts data source 336), and can provide data associated with potential attacks for further processing. For example, the alerts 360 may be indicative of potential attacks on an information technology network (e.g., the information technology network domain 102, shown in FIG. 1) and the alerts 362 may be indicative of potential attacks on an operational technology network (e.g., the operational technology network domain 104, shown in FIG. 1). Data maintained by the filtered alerts data source

336, for example, may be used for generating future security policies, as is discussed in further detail below in association with FIG. 5A.

[0027] FIG. 4A depicts an example system 400 that can execute implementations of the present disclosure. In the present example, the system 400 includes an alert aggregation system 402 (e.g., corresponding to the aggregator 214, shown in FIG. 2). The alert aggregation system 402, for example, can include various hardware and/or software-based components (e.g., software modules, objections, engines, libraries, etc.) including an alert fuser 410 and an alert aggregator 412. In general, the system 400 and its various components (e.g., components 410 and 412) can perform functions for processing and aggregating event/alert data received from various different sources. By aggregating event/alert data, for example, data redundancy can be decreased, and the aggregated event/alert data may be further processed to identify trends and correlations in the data.

[0028] In the present example, the aggregation system 402 can receive alert data 420 corresponding to potential attacks on an information technology network (e.g., the information technology network domain 102, shown in FIG. 1) and alert data 422 corresponding to potential attacks on an operational technology network (e.g., the operational technology network domain 104, shown in FIG. 1). Upon receiving the alert data 420 and the alert data 422, for example, the aggregation system 402 can use the fuser 410 to combine similar alerts, which may have been generated by different intrusion detection systems, security tools and/or sensors. For example, if multiple intrusion detection systems are included in the system 100, a malicious packet may be detected by each of the intrusion detection systems, and each of the systems may generate a similar alert in response to detecting the packet. In the present example, each of the similar alerts may include similar data, yet may have slightly different timestamps (e.g., due to network traffic speeds). If the fuser 410 determines that multiple alerts are related (e.g., the alerts were generated in response to the same packet or event based on having similar data and having timestamps within a threshold similarity value), for example, the multiple alerts may be combined into a meta-alert. The fuser 410 can provide meta-alerts 430 and raw alerts 432 (i.e., uncombined alerts) to the alert aggregator 412.

[0029] Upon receiving the meta-alerts 430 and raw alerts 432 from the alert fuser 410, for example, the alert aggregator 412 can aggregate the sets of alerts 430 and 432, based on data similarities. In general, alert aggregation may include combining alerts that have similar characteristics, which may indicate launch from and/or targeting of one or more computing devices. For example, an attack may include the scanning of particular computing devices included in an information technology network and computing devices included in an operational technology network, by multiple attackers. In the present example, alerts from the information technology network and alerts from the operational technology network may

be aggregated to reflect that the alerts are associated with the same type of attack.

[0030] In some implementations, each of the sets of alerts 430 and 432 may have similar data formats (e.g., an intrusion detection message exchange format (ID-MEF)), and may include data fields for source address, destination address, port number, timestamp, priority, and attack description. If the alert aggregator 412 determines that two or more of the meta-alerts 430 and/or the raw alerts 432 are similar based on criteria that pertains to data included in each alert (e.g., two or more alerts have the same destination address (and optionally, port number) and have timestamps within a threshold similarity value, two or more alerts have the same source address and have timestamps within a threshold similarity value, two or more alerts are targeted to similar services, or another suitable criteria based on similarity of alert attributes), for example, the alert aggregator 412 may aggregate the alerts. While the description above with regard to two or more of the meta-alerts 430 or raw alerts 432 being similar based on criteria that pertains to data included in each alert (and then being aggregated), it should be understood that the alert aggregator 412 may determine that a large or very large number of alerts are related (e.g. hundreds, thousands, or more alerts) and aggregate those alerts into a single record or data structure that facilitates efficient processing of a large volume of alert data. The alert aggregation system 402 can provide aggregated alerts 440 and non-aggregated alerts 442 as a combined set of meta and raw alerts 450 for further processing. The alert aggregation system 402 can also generate and store a data structure that represents both aggregated alerts and a relationship between different aggregated alerts, for example, based on a source and destination address in the aggregated alert. Once generated and stored, the data structure (described in greater detail below) can be further processed by the system, such as event correlation system 150 and system 500. In some implementations, the aggregated alerts 440 can be provided by the alert aggregation system 402 to an event analysis and management system 460 for further analysis.

[0031] FIG. 4B depicts an example data structure 480 that can be used by implementations of the present disclosure once it is generated and stored by the alert aggregation system 402. In some implementations, the data structure 480 (e.g., a directed graph) can be generated and used to represent the meta and raw alerts 450 provided by the alert aggregation system 402. Nodes (e.g., node 482, node 484) in the data structure 480 can represent entities referenced in the meta and raw alerts 450, such as computing devices/assets (e.g., devices/assets within the information technology (IT) network domain 102, the IT DMZ 106a, the operational technology (OT) network domain 104, and the OT DMZ 106b), internet protocol (IP) addresses, and ports. Edges (e.g., edge 486) in the data structure 480 can represent relationships between the entities (e.g., an aggregated event occurring

between two entities). For example, the edge 486 can represent a communication event (e.g., including a timestamp) between an originating computing device (e.g., a server) represented by the node 482 and a destination computing device (e.g., a DMZ server) represented by the node 484, and can be directed from the node 482 to the node 484. When representing computing devices and assets, addresses, and ports with unique nodes in a directed graph, for example, replication of entities may be avoided. When representing relationships between entities as edges in a directed graph, for example, multiple related alerts (e.g., alerts that include similar timestamp data and/or related message data) can be represented by a single edge between two nodes. For example, a particular computing device may attempt a series of port scans on each of a set of other devices. In the present example, the particular computing device (i.e., an origin) and each of the set of other devices (i.e., destinations) can be represented by nodes, and the port scans on each of the set of other devices can be consolidated and represented by a respective edge (e.g., an aggregated event or meta alert) between the particular computing device and each of the other devices.

[0032] In some implementations, when aggregating alert data and representing the alert data in a data structure, an edge that represents a relationship between two entities may be annotated with information that indicates a range of values included in a set of alerts that correspond to the edge. For example, if an originating device performs a series of port scans on a destination device over the course of a period of time, a single edge between a node representing the originating device and a node representing the destination device can be annotated with a value indicating a time of a first port scan during the period of time, a value indicating a time of a last port scan during the period of time, and a count of a total number of port scans during the period of time. As another example, if an originating device scans a range of ports on a destination device, a single edge between a node representing the originating device and a node representing the destination device can be appended with values indicating each of the ports scanned.

[0033] FIG. 5A depicts an example system 500 that can execute implementations of the present disclosure. In the present example, the system 500 includes an alert correlation and pattern extraction system 502 (e.g., corresponding to the correlator 216 and the pattern recognizer 218, shown in FIG. 2). The alert correlation and pattern extraction system 502 can include various hardware and/or software-based components (e.g., software modules, objects, engines, libraries, etc.) including an alert extractor 510, an alert correlator 512, a pattern extractor 514, a threat analyzer 516, and a visualization generator 518. Various data sources (e.g., databases, file systems, etc.) may maintain data used by the system 500 and its components. In the present example, the system 500 includes a targets and attackers data source 530, a threat intelligence data source 534 (e.g., similar

to the threat intelligence data source 234, shown in FIG. 2), and an unsuccessful attacks data source 536 (e.g., including data provided by the filtered alerts data source 336, shown in FIG. 3). In general, the system 500 and its various components (e.g., components 510, 512, 514, 516, and 518) can perform functions for processing and correlating aggregated event/alert data. By correlating aggregated alert data, for example, complex multi-step attacks against an entire industrial control system network (e.g., including information technology and operational technology network domains) may be detected.

[0034] In the present example, the correlation and pattern extraction system 502 can receive meta and raw alerts 550 (e.g., similar to the meta and raw alerts 450, shown in FIG. 4A) from the alert aggregation system 402 (shown in FIG. 4A). The meta and raw alerts 550, for example, may include aggregated and non-aggregated alerts that are associated with suspicious network activity from multiple different network domains (e.g., the industrial technology (IT) network domain 102, the IT DMZ 106a, the operational technology (OT) network domain 104, and the OT DMZ 106b, shown in FIG. 1). In some implementations, the meta and raw alerts may 550 may be represented by a data structure (e.g., the data structure 480, shown in FIG. 4B). Upon receiving the meta and raw alerts 550, for example, the correlation and pattern extraction system 502 can provide alert data to the alert extractor 510 and to the alert correlator 512. The alert extractor 510, for example, can analyze the meta and raw alerts 550, can extract information from the meta and raw alerts, and can generate a list of computing devices that are targets of attacks, along with data corresponding to attackers of the targets, such as internet protocol (IP) addresses of the attackers. In some implementations, the alert extractor 510 can traverse the data structure 480 (e.g., a directed graph) and mine the data structure for internet protocol (IP) addresses, uniform resource locator (URL) information, and other relevant data. Data associated with targets and attackers can be provided by the alert extractor 510 to the targets and attackers data source 530, for example. The alert correlator 512, for example, can generate one or more threat scenarios 560 (e.g., chains of events/alerts that indicate attack paths), based on data (e.g., attack signatures and profiles) provided by the threat intelligence data source 534.

[0035] In general, correlating alert data may include enriching (e.g., labeling) the alert data, linking related alerts, and identifying an attack path indicated by the linked alerts. For example, the alert correlator 512 can generate threat scenarios 560 by determining chains of alerts included in the meta and raw alerts 550, and optionally prioritizing, labeling, or otherwise enriching alert data, based on data provided by the threat intelligence data source 534. Attack signatures and profiles within the threat intelligence data source 534, for example, may relate to communication patterns between computing devices, and may include information related to potential

target devices for a type of attack. As new attacks are detected, for example, information related to attack signatures and profiles may be added to the threat intelligence data source 534.

[0036] Enriching alert data, for example, may include analyzing alert data (e.g., the meta and raw alerts 550) and identifying alerts that are associated with a particular attacker. For example, a multi-stage attack performed by an attacker may include reconnaissance, delivery, and installation stages. Each stage, for example, may include communication between the attacker and one or more computing devices on a network during one or more sessions. When enriching the meta and raw alerts 550, for example, the alert correlator 512 can use information provided by the targets and attackers data source 530, the threat intelligence data source 534, and/or additional threat intelligence services to identify alerts that are associated with communications from the particular attacker (e.g., indicated by Internet Protocol (IP) address and country of origin information), and can label or otherwise enrich information related to the alerts as being potentially related to a multi-stage attack.

[0037] Linking alert data, for example, may include analyzing previously enriched (e.g., labeled) alerts/events. For example, a subset of the meta and raw alerts 550 may be labeled as being related to an attack performed by a previously identified attacker. The subset of alerts, for example, may correspond to communication with the attacker and with subsequent alerts that may result from the communication (e.g., after the delivery of a malicious payload). In some implementations, a time-based threshold may be used to determine whether two or more alerts are to be linked. For example, the alert correlator 512 may link two or more similarly labeled alerts if the alerts have timestamps within a suitable time threshold value (e.g., one minute, five minutes, ten minutes, or another suitable value). The time threshold value, for example, may be a configurable tuning parameter. In some implementations, linking alert data (i.e., determining a chain of alerts) may depend on a type of attack indicated by an attack signature and/or profile, or another sort of indicator (e.g., an address of an attacker and/or target). For example, a different sub-function may be used by the alert correlator 512 for determining a chain of alerts for each different type of attack, based on attack signatures and profiles defined by the threat intelligence data source 534.

[0038] Identifying an attack path indicated by the linked alerts, for example, may include identifying steps that an attacker takes toward an intended target. The alert correlator 512, for example, can analyze previously linked alerts to identify an attack path that may cross one or more network domains. In general, identifying an attack path may include determining a series of communications between computing devices. For example, an attacker may attempt to change the status of a controller device in an operational technology (OT) network, but lack physical access. In the present example, to gain

access to the controller device, the attacker may launch an attack against a computer device in an information technology (IT) network, and may exploit the computer device in order to step to a human-machine interface (HMI) device in the operational technology network, and then may further exploit the human-machine interface device in order to step to the controller device. Information associated with identified attack paths can be used by the alert correlator 512, to determine one or more threat scenarios 560.

[0039] In some implementations, correlating alert data may include traversing a data structure and determining whether paths included in the data structure are of interest (e.g., represent an attack) or are irrelevant (e.g., represent false positives, noise, etc.). Referring to FIG. 5B, for example, an example data structure 580 (e.g., a directed graph) is shown. The alert correlator 512, for example, can use information provided by the threat intelligence data source 534 and/or other rule or pattern-based sources and/or anomaly detection mechanisms to determine which paths may be of interest, and which paths may be irrelevant. As another example, an anomaly detection mechanism may observe an unusual occurrence of communication between computing devices (e.g., communication between an external host and a programmable logic controller (PLC)), an unusual rate of events between computing devices in a particular time-frame, or another sort of anomaly. Paths including nodes and edges corresponding to one or more anomalous occurrences, for example, may be of interest.

[0040] In some implementations, paths that are determined to be irrelevant may be pruned from a data structure. In the present example, nodes 482a, 482b, 482c, and 482d can represent computing devices (e.g., servers) within the information technology (IT) network domain 102, the IT DMZ 106a, the operational technology (OT) network domain 104, or the OT DMZ 106b. The alert correlator 512, for example, may determine that the nodes 482a, 482b, 482c, and 482d are without relationship edges directed to them, and that the nodes do not represent originating devices (i.e., devices that are identified as originating an attack path). Thus, in the present example, the nodes 482a, 482b, 482c, and 482d may be determined as being irrelevant to a path of interest, and the nodes and any edges directed from the nodes that represent relationships to other computing devices represented in the data structure 580 may be pruned by the alert correlator 512. The pruned data structure can then be stored for further processing and/or for display to a system operator.

[0041] In some implementations, information related to paths that have previously been identified (e.g., by an operator) as being irrelevant may be used to prune paths from a data structure. For example, a path that had previously been determined by the alert correlator 512 as being a potential path of interest may be identified by an operator as being irrelevant (e.g., a false positive). In the present example, if the alert correlator 512 determines

that a particular path includes characteristics of the path that was identified as being irrelevant (e.g., based on information provided by the threat intelligence data source 534), the particular path may be pruned from the data structure 580.

[0042] Referring to FIG. 5C, for example, an example data structure 590 (e.g., a pruned version of the example data structure 580) is shown. The alert correlator 512 can use information from the threat intelligence data source 534, for example, to enrich the example data structure 590 to include additional information regarding one or more paths of interest, such as information related to past actions associated with computing devices represented by nodes included in the data structure. For example, if a particular computing device has been reported as being associated with malicious activities, a node representing the computing device may be enriched to include such information.

[0043] Referring again to FIG. 5A, the threat scenarios 560 can be provided by the event correlator 512, for example, to the threat analyzer 516 and the visualization generator 518. In some implementations, the threat scenarios 560 may be represented by a pruned data structure. For example, the data structure 590 (e.g., a pruned directed graph, shown in FIG. 5C) may include various paths of interest (i.e., potential attack paths), enriched with additional information. In some implementations, the alert correlator 512 can provide uncorrelated meta-alerts 562 to an event analysis and management system 570 (e.g., a security information and event management (SIEM) system). For example, an uncorrelated meta-alert may eventually be correlated to other alerts (e.g., after additional event/alert data is collected), and the previously uncorrelated meta-alert may be correlated to generate an additional threat scenario. In some implementations, the alert correlator 512 can provide correlated meta-alerts 562 to the event analysis and management system 570. For example, the event analysis and management system 570 may support a format to receive correlated and/or uncorrelated meta-alerts.

[0044] The alert correlation and pattern extraction system 502 can use the pattern extractor 514, for example, to detect previously unknown security threats. For example, the pattern extractor 514 can analyze data from the targets and attackers data source 530 and from the unsuccessful attacks data source 536 to identify patterns, signatures, and rules associated with potential security attacks. In some implementations, data related to potential attacks can be provided to one or more components for use in identifying and/or preventing future attacks. In the present example, data from the pattern extractor 514 can be provided to a risk management component 572 and/or to the threat intelligence data source 534. In general, the risk management component 572 may provide information to a user (e.g., a network administrator) to assist in installing new software or software patches within a system, based on identified risks.

[0045] The alert correlation and pattern extraction sys-

tem 502 can use the threat analyzer 516, for example, to analyze the threat scenarios 560, and to generate data associated with one or more courses of action 574. For example, the threat analyzer 516 can determine the impact of attacks on the operation of information technology and operational technology networks (e.g., the information technology network domain 102 and the operational technology network domain 104, shown in FIG. 1), rank threat scenarios 560 based on their importance and on system priorities, and provide courses of action 574 relevant to each threat scenario (e.g., based on information provided by the threat intelligence data source 534). Processed threat scenario data, for example, can be provided by the threat analyzer 516 to the threat intelligence data source 534. In some implementations, appropriate courses of action 574 may be provided for each domain (e.g., information technology and operational technology) and/or each device in an industrial control system network. Courses of action 574, for example, may include actions such as closing ports on computing devices, blocking communications that originate from particular internet protocol addresses, shutting down computing devices, and so forth.

[0046] The alert correlation and pattern extraction system 502 can use the visualization generator 518, for example, to generate one or more reports and/or diagrams (e.g., diagram 576). For example, the visualization generator 518 can analyze the threat scenarios 560 and can determine a progression of system states and communication paths during an attack. In general, diagrams generated by the visual generator 518 may include sunburst diagrams, node diagrams, and/or other suitable diagrams, which may illustrate communications patterns between computing devices in a network. In some implementations, a visualization may be generated based on a data structure that represents one or more potential paths of interest (e.g., attack paths). For example, the data structure 590 (e.g., a pruned directed graph, shown in FIG. 5C) can be used by the visualization generator 518 to generate the diagram 576. The diagram 576, for example, can be presented to an operator by the interface device 154 (shown in FIG. 1). Upon reviewing the diagram 576, for example, the operator may identify anomalous network activity and/or may identify one or more irrelevant paths (e.g., false positives, noise, etc.).

[0047] FIG. 6 is a flowchart of an example process 600 that can be executed in accordance with implementations of the present disclosure. The process 600, for example, can be performed by systems such as one or more of the example systems described above. Briefly, the example process 600 includes receiving activity data from multiple domains, filtering and verifying the activity data, aggregating the activity data, correlating the activity data, and providing one or more visualizations and courses of action in response to the activity data.

[0048] Activity data can be received from multiple domains (602). Referring to FIGS. 1 and 2 and as discussed above, for example, activity data (e.g., event/alert data

provided by one or more intrusion detection systems) can be received by an event correlation system. In the present example, activity data can be received from an information technology (IT) and from an operational technology (OT) network.

[0049] Activity data can be filtered and verified (604). Referring to FIGS. 2 and 3 and as discussed above, for example, information technology (IT) and operational technology (OT) activity data can be filtered, verified, and further processed. Filtered activity data can be maintained for further use (e.g., by one or more threat intelligence services). Activity data (e.g., raw alerts) indicative of potential attacks can be provided for aggregation.

[0050] Activity data can be aggregated (606). Referring to FIGS. 2 and 4 and as discussed above, for example, raw alerts from an information technology (IT) and from an operational technology (OT) network can be fused and aggregated. Fused and aggregated alerts (i.e., meta-alerts) and raw alerts can be provided for correlation.

[0051] Activity data can be correlated (608). Referring to FIGS. 2 and 5 and as discussed above, for example, meta-alerts and raw alerts can be correlated, based on threat intelligence (e.g., attack signatures and profiles). Further, the meta-alerts and raw alerts can be extracted to identify data associated with targets and attackers. Based on data associated with the targets and attackers and based on previously filtered activity data, for example, patterns can be extracted and the threat intelligence can be updated.

[0052] One or more visualizations and/or courses of action can be provided (610). Referring to FIGS. 2 and 5 and as discussed above, for example, by correlating meta-alerts and raw alerts, one or more threat scenarios may be identified, which can be used for generating visualizations and/or for identifying courses of action. Further, information associated with meta-alerts may be provided to an event analysis and management system.

[0053] Embodiments of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions encoded on a tangible non-transitory program carrier for execution by, or to control the operation of, data processing apparatus. Alternatively or in addition, the program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. The computer storage medium can be a machine-readable storage device, a machine-readable storage sub-

strate, a random or serial access memory device, or a combination of one or more of them.

[0054] The term "data processing apparatus" refers to data processing hardware and encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can also be or further include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The apparatus can optionally include, in addition to hardware, code that creates an execution environment for computer programs, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[0055] A computer program, which may also be referred to or described as a program, software, a software application, a module, a software module, a script, or code, can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub-programs, or portions of code. A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

[0056] The processes and logic flows described in this specification can be performed by one or more programmable computers executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit).

[0057] Computers suitable for the execution of a computer program include, by way of example, general or special purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read-only memory or a random access memory or both. The essential elements of a computer are a central processing unit for performing or executing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto-optical disks, or optical

disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

[0058] Computer-readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0059] To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's device in response to requests received from the web browser.

[0060] Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front-end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back-end, middleware, or front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network (LAN) and a wide area network (WAN), e.g., the Internet.

[0061] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other. In some embodiments, a server transmits data, e.g., an HTML page, to a user device, e.g., for purposes

of displaying data to and receiving user input from a user interacting with the user device, which acts as a client. Data generated at the user device, e.g., a result of the user interaction, can be received from the user device at the server.

[0062] An example of one such type of computer is shown in FIG. 7, which shows a schematic diagram of a generic computer system 700. The system 700 can be used for the operations described in association with any of the computer-implement methods described previously, according to one implementation. The system 700 includes a processor 710, a memory 720, a storage device 730, and an input/output device 740. Each of the components 710, 720, 730, and 740 are interconnected using a system bus 750. The processor 710 is capable of processing instructions for execution within the system 700. In one implementation, the processor 710 is a single-threaded processor. In another implementation, the processor 710 is a multi-threaded processor. The processor 710 is capable of processing instructions stored in the memory 720 or on the storage device 730 to display graphical information for a user interface on the input/output device 740.

[0063] The memory 720 stores information within the system 700. In one implementation, the memory 720 is a computer-readable medium. In one implementation, the memory 720 is a volatile memory unit. In another implementation, the memory 720 is a non-volatile memory unit.

[0064] The storage device 730 is capable of providing mass storage for the system 700. In one implementation, the storage device 730 is a computer-readable medium. In various different implementations, the storage device 730 may be a floppy disk device, a hard disk device, an optical disk device, or a tape device.

[0065] The input/output device 740 provides input/output operations for the system 700. In one implementation, the input/output device 740 includes a keyboard and/or pointing device. In another implementation, the input/output device 740 includes a display unit for displaying graphical user interfaces.

[0066] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or on the scope of what may be claimed, but rather as descriptions of features that may be specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombina-

tion or variation of a subcombination.

[0067] Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0068] Particular implementations of the subject matter have been described. Other implementations are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In some cases, multitasking and parallel processing may be advantageous.

Embodiments

[0069] Although the present invention is defined in the attached embodiments, it should be understood that the present invention can also (alternatively) be defined in accordance with the following groups of embodiments:

Group 1:

1. A computer-implemented method for correlating domain activity data, the method comprising:
 - receiving first domain activity data from a first network domain and second domain activity data from a second network domain, the first domain activity data and the second domain activity data including events, alerts, or both from the respective first and second network domains;
 - filtering the first domain activity data and the second domain activity data to remove irrelevant activity data, based on a first set of profile data for devices in the first network domain and a second set of profile data for devices in the second network domain;
 - aggregating unfiltered first domain activity data and unfiltered second domain activity data;
 - correlating aggregated unfiltered first domain activity data and unfiltered second domain activity data to determine an attack

path for an attack that occurs across the first network domain and the second network domain, based on attack signatures and profiles associated with previously identified attacks; and
generating a visualization of the attack path.

2. The method of embodiment 1, wherein the first domain activity data, the second domain activity data, or both, includes log data provided by one or more security sensors.

3. The method of embodiment 1, wherein the first network domain is an information technology network domain and the second network domain is an operational technology network domain.

4. The method of embodiment 1, wherein the filtering comprises:

- determining, for each event or alert, a corresponding attack and a corresponding target;
- determining, based on profile data for the corresponding target, that the attack on the target is rendered unsuccessful; and
- filtering the corresponding event or alert.

5. The method of embodiment 4, wherein the filtering comprises:

- for each unfiltered event or alert, dynamically retrieving current status information about the corresponding target;
- determining, based the current status information about the corresponding target, that the attack on the target is rendered unsuccessful; and
- filtering the corresponding event or alert.

6. The method of embodiment 1, wherein the aggregating comprises determining that two or more alerts were generated in response to detecting a same packet, and combining the alerts into a meta-alert.

7. The method of embodiment 6, wherein the aggregating comprises determining that two or more alerts, meta-alerts, or both, are associated with a same attack or have similar characteristics, and combining the alerts, meta-alerts, or both.

8. The method of embodiment 7, wherein the aggregating is performed when the two or more alerts, meta-alerts, or both have timestamps within a threshold similarity value and include

the same destination address, the same source address, or both.

9. The method of embodiment 1, wherein the correlating comprises:

labeling the aggregated unfiltered first domain activity data and unfiltered second domain activity data to identify two or more alerts, meta-alerts, or both that are associated with a particular attacker; linking the activity data that is labeled as being associated with the particular attacker to identify a chain of two or more alerts, meta-alerts, or both; and determining the attack path that occurs across the first network domain and the second network domain, including determining a series of communications between one or more devices in the first network domain and one or more devices in the second network domain.

10. The method of embodiment 1, wherein two or more alerts, meta-alerts, or both are linked when the alerts or meta-alerts have timestamp values within a time threshold value.

11. The method of embodiment 1, further comprising:

based on filtered first domain activity data and filtered second domain activity data, determining and storing filtered data associated with unsuccessful attacks; and based on aggregated unfiltered first domain activity data and unfiltered second domain activity data, determining and storing data associated with targets and attackers; wherein the attack signatures and profiles are based on the filtered data associated with unsuccessful attacks and on the data associated with targets and attackers.

12. The method of embodiment 11, wherein the data associated with targets and attackers includes references to devices that are targets of attacks, and references to addresses of attackers.

13. The method of embodiment 1, further comprising:

receiving data associated with the attack path; in response to receiving the data associated with the attack path, determining an impact of the attack on the first network domain and

on the second network domain; and providing an appropriate course of action for the first network domain and the second network domain.

14. The method of embodiment 1, further comprising:

correlating aggregated unfiltered first domain activity data and unfiltered second domain activity data to determine multiple attack paths; for each of the attack paths, determining an impact of the attack on the first network domain and on the second network domain; and ranking each of the multiple attack paths, based on the impact of the respective attack.

Group 2:

1. A computer-implemented method for correlating domain activity data, the method comprising:

receiving first domain activity data from a first network domain and second domain activity data from a second network domain, the first domain activity data and the second domain activity data including events, alerts, or both from the respective first and second network domains; filtering the first domain activity data and the second domain activity data to remove irrelevant activity data, based on a first set of profile data for devices in the first network domain and a second set of profile data for devices in the second network domain; aggregating unfiltered first domain activity data and unfiltered second domain activity data; correlating aggregated unfiltered first domain activity data and unfiltered second domain activity data to determine an attack path for an attack that occurs across the first network domain and the second network domain, based on attack signatures and profiles associated with previously identified attacks; and generating a visualization of the attack path.

2. The method of embodiment 1, wherein the first domain activity data, the second domain activity data, or both, includes log data provided by one or more security sensors.

3. The method of embodiment 1, wherein the

first network domain is an information technology network domain and the second network domain is an operational technology network domain.

5

4. The method of embodiment 1, wherein the filtering comprises:

determining, for each event or alert, a corresponding attack and a corresponding target;

10

determining, based on profile data for the corresponding target, that the attack on the target is rendered unsuccessful; and

filtering the corresponding event or alert.

15

5. The method of embodiment 4, wherein the filtering comprises:

for each unfiltered event or alert, dynamically retrieving current status information about the corresponding target;

20

determining, based the current status information about the corresponding target, that the attack on the target is rendered unsuccessful; and

25

filtering the corresponding event or alert.

6. The method of embodiment 1, wherein the aggregating comprises determining that two or more alerts were generated in response to detecting a same packet, and combining the alerts into a meta-alert.

30

7. The method of embodiment 6, wherein the aggregating comprises determining that two or more alerts, meta-alerts, or both, are associated with a same attack or have similar characteristics, and combining the alerts, meta-alerts, or both.

35

40

8. The method of embodiment 7, wherein the aggregating is performed when the two or more alerts, meta-alerts, or both have timestamps within a threshold similarity value and include the same destination address, the same source address, or both.

45

9. The method of embodiment 1, wherein the correlating comprises:

50

labeling the aggregated unfiltered first domain activity data and unfiltered second domain activity data to identify two or more alerts, meta-alerts, or both that are associated with a particular attacker;

55

linking the activity data that is labeled as being associated with the particular attacker

to identify a chain of two or more alerts, meta-alerts, or both; and

determining the attack path that occurs across the first network domain and the second network domain, including determining a series of communications between one or more devices in the first network domain and one or more devices in the second network domain.

10. The method of embodiment 1, wherein two or more alerts, meta-alerts, or both are linked when the alerts or meta-alerts have timestamp values within a time threshold value.

11. The method of embodiment 1, further comprising:

based on filtered first domain activity data and filtered second domain activity data, determining and storing filtered data associated with unsuccessful attacks; and

based on aggregated unfiltered first domain activity data and unfiltered second domain activity data, determining and storing data associated with targets and attackers; wherein the attack signatures and profiles are based on the filtered data associated with unsuccessful attacks and on the data associated with targets and attackers.

12. The method of embodiment 11, wherein the data associated with targets and attackers includes references to devices that are targets of attacks, and references to addresses of attackers.

13. The method of embodiment 1, further comprising:

receiving data associated with the attack path;

in response to receiving the data associated with the attack path, determining an impact of the attack on the first network domain and on the second network domain; and providing an appropriate course of action for the first network domain and the second network domain.

14. The method of embodiment 1, further comprising:

correlating aggregated unfiltered first domain activity data and unfiltered second domain activity data to determine multiple attack paths; for each of the attack paths, determining an

impact of the attack on the first network domain and on the second network domain;
and
ranking each of the multiple attack paths, based on the impact of the respective attack.

Claims

1. A computer-implemented method for transforming representations of network activity data, the method being executed by one or more processors and comprising:

receiving a data structure that represents communication events between computing devices of one or more networks, wherein the data structure is a directed graph stored in a graph database;
analyzing the data structure and determining a set of potential attack paths represented in the data structure;
assigning a score to each potential attack path in the set of potential attack paths;
removing potential attack paths from the set of potential attack paths that have scores that do not meet a predetermined threshold;
ranking potential attack paths that remain in the set of potential attack paths, based on each score assigned to each potential attack path;
and
providing the data structure that includes a ranked set of potential attack paths.

2. The method of claim 1, wherein the computing devices of the one or more networks are represented by nodes of the directed graph, and the communication events are represented by directed edges between the nodes of the directed graph.
3. The method of claim 2, wherein determining the set of potential attack paths represented in the data structure includes traversing the directed graph based on timestamp information associated with the directed edges between the nodes to determine a plurality of node sequences, each node sequence proceeding from an originating node to a destination node.
4. The method of claim 3, wherein determining the set of potential attack paths represented in the data structure includes determining all possible node sequences of two or more linked nodes within the data structure.
5. The method of claim 1, wherein analyzing the data structure includes receiving information provided by

a rule or pattern-based threat intelligence data source.

6. The method of claim 1, wherein analyzing the data structure includes determining a meshedness coefficient for the data structure that indicates whether one or more potential attack paths represented in the data structure are looped.
7. The method of claim 6, wherein determining the set of potential attack paths represented in the data structure includes converting looped potential attack paths to non-looped potential attack paths.
8. The method of claim 1, wherein assigning the score to each potential attack path in the set of potential attack paths includes determining, for each potential attack path, two or more component scores including two or more of a spatial component score, a temporal component score, and an importance component score, and wherein the score for the potential attack path is an aggregation of the two or more component scores.
9. The method of claim 8, wherein the spatial component score represents a distance between computing devices within the one or more networks.
10. The method of claim 8, wherein the spatial component score represents a number of logical system boundaries crossed by the potential attack path.
11. The method of claim 8, wherein the temporal component score represents a function of the rate at which communication events associated with the attack path occur.
12. The method of claim 8, wherein the importance component score represents an importance of one or more computing devices associated with the potential attack path, and optionally wherein determining the importance component score includes receiving information related to the one or more computing devices associated with the potential attack path from a configuration management system.
13. The method of claim 8, wherein the importance component score represents an importance of one or more communication events associated with the potential attack path, and optionally wherein determining the importance component score includes receiving information related to the one or more communication events associated with the potential attack path from a network security sensor.
14. A system, comprising:

one or more processors; and

a computer-readable storage device coupled to the one or more processors and having instructions stored thereon which, when executed by the one or more processors, cause the one or more processors to perform operations for transforming representations of network activity data, the operations comprising the method of any one of claims 1 to 13. 5

15. A non-transitory computer-readable storage medium coupled to one or more processors and having instructions stored thereon which, when executed by the one or more processors, cause the one or more processors to perform operations for transforming representations of network activity data, the operations comprising the method of any one of claims 1 to 13. 10 15

20

25

30

35

40

45

50

55




FIG. 1




FIG. 2




FIG. 3




FIG. 4A




FIG. 4B






FIG. 5A



580 ↗

FIG. 5B



590

FIG. 5C




FIG. 5D




FIG. 5E

600 ↘





FIG. 6A

620 ↘

**FIG. 6B**




FIG. 7



EUROPEAN SEARCH REPORT

Application Number
EP 16 16 4616

5

10

15

20

25

30

35

40

45

50

55

1

EPO FORM 1503 03.82 (P04C01)

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	US 2011/252032 A1 (FITZGERALD ROBERT E [US] ET AL) 13 October 2011 (2011-10-13) * paragraph [0004] * * paragraph [0031] * * paragraph [0040] * * paragraph [0054] * * paragraph [0063] * * paragraphs [0087] - [0091] * * paragraph [0132] * * figures 14-15 *	1-15	INV. H04L29/06
X	US 2007/209074 A1 (COFFMAN THAYNE R [US]) 6 September 2007 (2007-09-06) * paragraphs [0012] - [0015] * * paragraphs [0020] - [0022] * * paragraph [0097] * * paragraphs [0142] - [0145] * * figure 4D * * figures 5A-5B *	1-15	
A	US 8 307 444 B1 (MAYER ALAIN JULES [US] ET AL) 6 November 2012 (2012-11-06) * column 7, lines 10-20 * * column 8, lines 45-60 * * column 9, lines 10-25 * * column 10, lines 40-50 * * column 11, lines 55-65 * * column 12, lines 5-30 * * figure 4A *	1-15	TECHNICAL FIELDS SEARCHED (IPC) H04L H04W G06F
A	US 2015/074806 A1 (ROUNDY KEVIN [US] ET AL) 12 March 2015 (2015-03-12) * paragraph [0036] * * paragraph [0039] * * paragraph [0056] * * paragraph [0065] * * paragraphs [0070] - [0073] * * figure 3 *	1-15	
The present search report has been drawn up for all claims			
Place of search Munich		Date of completion of the search 2 June 2016	Examiner Erdene-Ochir, O
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 16 16 4616

5

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

02-06-2016

10

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011252032 A1	13-10-2011	NONE	
US 2007209074 A1	06-09-2007	US 2007209074 A1 WO 2008021585 A2	06-09-2007 21-02-2008
US 8307444 B1	06-11-2012	US 7890869 B1 US 8132260 B1 US 8307444 B1 US 8321944 B1	15-02-2011 06-03-2012 06-11-2012 27-11-2012
US 2015074806 A1	12-03-2015	US 2015074806 A1 WO 2015035559 A1	12-03-2015 19-03-2015

15

20

25

30

35

40

45

50

55

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 62145248 A [0001]